arXiv:1402.6536v1 [stat.ME] 26 Feb 2014

On split sample and randomized confidence intervals for binomial proportions

Måns Thulin

Department of Mathematics, Uppsala University

Abstract

Split sample methods have recently been put forward as a way to reduce the coverage oscillations that haunt confidence intervals for parameters of lattice distributions, such as the binomial and Poisson distributions. We study split sample intervals in the binomial setting, showing that these intervals can be viewed as being based on adding discrete random noise to the data. It is shown that they can be improved upon by using noise with a continuous distribution instead, regardless of whether the randomization is determined by the data or an external source of randomness. We compare split sample intervals to the randomized Stevens interval, which removes the coverage oscillations completely, and find the latter interval to have several advantages.

Keywords: Binomial distribution; confidence interval; Poisson distribution; randomization; split sample method.

1 Introduction

The problem of constructing confidence intervals for parameters of discrete distributions continues to attract considerable interest in the statistical community. The lack of smoothness of these distributions causes the coverage (i.e. the probability that the interval covers the true parameter value) of such intervals to fluctuate from $1-\alpha$ when either the parameter value or the sample size n is altered. Recent contributions to the theory of these confidence intervals include Krishnamoorthy & Peng (2011), Newcombe (2011), Gonçalves et al. (2012), Göb & Lurz (2013) and Thulin (2013a).

The purpose of this short note is to discuss the split sample method recently proposed by Decrouez & Hall (2014). The split sample method reduces the oscillations of the coverage by splitting the sample in two. It is applicable to most existing confidence intervals for parameters of lattice distributions, including the

binomial and Poisson distributions. For simplicity, we will in most of the remainder of the paper limit our discussion to the binomial setting, with the split sample method being applied to the celebrated Wilson (1927) interval for the proportion p. Our conclusions are however equally valid for other confidence intervals and distributions.

A random variable $X \sim \text{Bin}(n, p)$ is the sum of a sequence $X_1, \ldots, X_n$ of independent Bernoulli(p) random variables. The split sample method is applied to $X_1, \ldots, X_n$ rather than to X. The idea is to split the sample into two sequences $X_1, \ldots, X_{n_1}$ and $X_{n_1+1}, \ldots, X_{n_1+n_2}$ with $n_1 + n_2 = n$ and $n_1 \neq n_2$. In the formula for the chosen confidence interval, the maximum likelihood estimator $\hat{p} = X/n$ is then replaced by the weighted estimator

$$\tilde{p} = \frac{1}{2}\Big(\frac{1}{n_1}\sum_{i=1}^{n_1} X_i + \frac{1}{n_2}\sum_{i=n_1+1}^{n_1+n_2} X_i\Big). \tag{1}$$

The formula for the Wilson interval is

$$\frac{\hat{p}n + z_{\alpha/2}^2/2}{n + z_{\alpha/2}^2} \pm \frac{z_{\alpha/2}}{n + z_{\alpha/2}^2}\sqrt{\hat{p}(1-\hat{p})n + z_{\alpha/2}^2/4},$$

where $z_{\alpha/2}$ is the $\alpha/2$ standard normal quantile, so the split sample Wilson interval is given by

$$\frac{\tilde{p}n + z_{\alpha/2}^2/2}{n + z_{\alpha/2}^2} \pm \frac{z_{\alpha/2}}{n + z_{\alpha/2}^2}\sqrt{\tilde{p}(1-\tilde{p})n + z_{\alpha/2}^2/4}.$$

Decrouez & Hall (2014) showed by numerical and asymptotic arguments based on Decrouez & Hall (2013) that when $n_1 = n/2 + 0.15n^{3/4}$ (rounded to the nearest integer) and $n_2 = n - n_1$ the split sample method greatly reduces the oscillations of the coverage of the interval, without increasing the interval length. In the remainder of the paper, whenever n_1 and n_2 need to be specified, we will use these values.

Depending on how the sample is split, different confidence intervals will result, making the split sample interval a randomized confidence interval. If the sequence $X_1, \ldots, X_n$ is available, the interval can be data-randomized in the sense that the randomization can be determined by the data: the first n_1 observations can be put in the first subsample and the remaining n_2 observation in the second subsample. If the results of the individual trials have not been recorded, one must use randomness from outside the data to create a sequence $X_1, \ldots, X_n$ of 0's and 1's such that $\sum_{i=1}^n X_i = X$. We will discuss these two settings separately.

Decrouez & Hall (2014) left two questions open. The first is how the split sample interval performs in comparison to other randomized intervals, as Decrouez & Hall (2014) only compared the split sample interval to non-randomized confidence

Figure 1: The distribution of $Y|X$ when $n = 11$, $n_1 = 6$ and $n_2 = 5$.

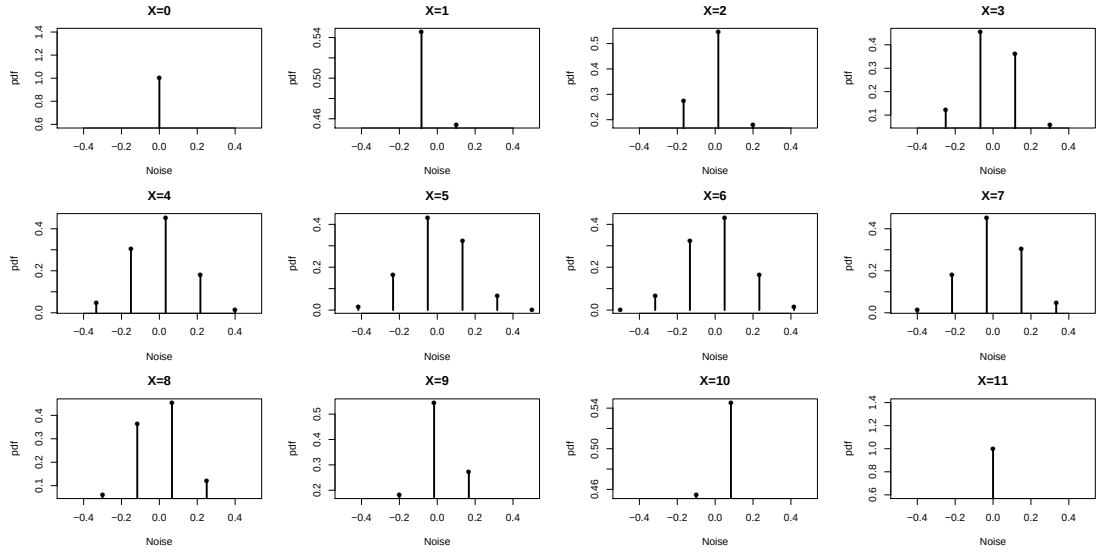

intervals. The second question is to what extent the randomization can affect the bounds of the interval. In the remainder of this note we answer these questions, discussing the impact of the random splitting on the confidence interval and comparing the split sample interval to alternative intervals. Section 2 is concerned with externally randomized intervals whereas we in Section 3 study data-randomized intervals. In both settings, we find the competing intervals to be superior to the split sample interval. Various criticisms of randomized intervals are then discussed in Section 4.

2 Externally randomized confidence intervals

2.1 Split sample and alternative intervals

A strategy for smoothing the distribution of the binomial random variable X is to base our inference on $X + Y$, where Y is a comparatively small random noise, using $X+Y$ instead of X in the formula for our chosen confidence interval. Having a smoother distribution leads to a better normal approximation, which in turn reduces the coverage fluctuations of the interval. The split sample method can be seen to be a special case of this strategy. Let Z be a random variable which, conditioned on X, follows a Hypergeometric(n, X, n_1) distribution. Then it follows

Figure 2: Comparison between the cumulative distribution functions of normal distributions (red) and (a) X, (b) $\tilde{X}$, (c) $\dot{X}$ when $X \sim \text{Bin}(7, 1/2)$.

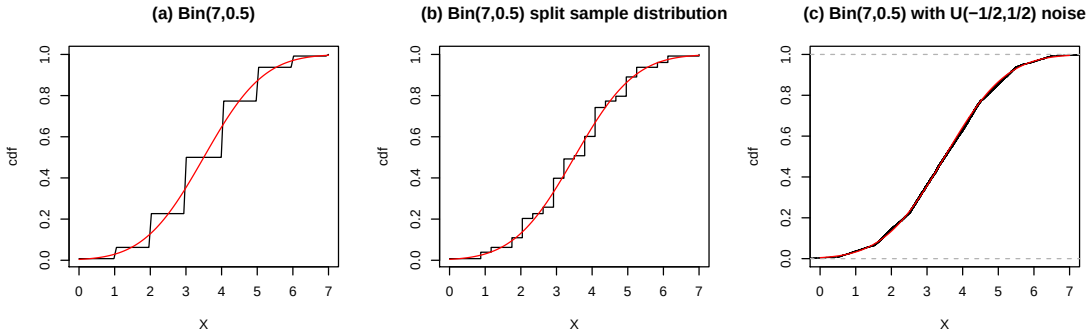

from (1) that

$$\tilde{X} = n\tilde{p} \stackrel{d}{=} \frac{n}{2n_1} Z - \frac{n}{2n_2}(X - Z),$$

so that

$$\tilde{X} \stackrel{d}{=} X + Y \qquad \text{with} \qquad Y = \frac{n}{2n_1} Z - \frac{n}{2n_2} Z + \frac{n_1 - n_2}{2n_2} X.$$

The conditional distribution of Y when $n = 11$ is shown in Figure 1. Indeed, when the sequence $X_1, \ldots, X_n$ has not been observed, the splitting of the sample is superficial (since the sample itself is not available to us) and it may be more natural to think of the split sample procedure as adding this discrete random noise term Y, the distribution of which is conditioned on X.

If one has decided to use random noise to improve the normal approximation, the next step is to ask whether there are other distributions for the noise that yield an even better approximation and thereby decrease the coverage oscillations even further. The answer is yes, for instance if X is replaced by $\dot{X} = X + Y$ where $Y \sim \text{U}(-1/2, 1/2)$ independently of X. The normal approximations of X, $\tilde{X}$ and $\dot{X}$ are compared in Figure 2. We will refer to the interval based on $\dot{X}$ as the $\text{U}(-1/2, 1/2)$ interval.

A randomized confidence interval that does not rely on a normal approximation is the Stevens (1950) interval. It belongs to an important class of confidence intervals for p, consisting of intervals (p_L, p_U) where the lower bound p_L is such that

$$\nu_1 \cdot \binom{n}{X} p_L^X (1 - p_L)^{n-X} + \sum_{k=X+1}^{n} \binom{n}{k} p_L^k (1 - p_L)^{n-k} = \alpha/2$$

Figure 3: Coverage and expected length of some confidence intervals for the binomial proportion p.

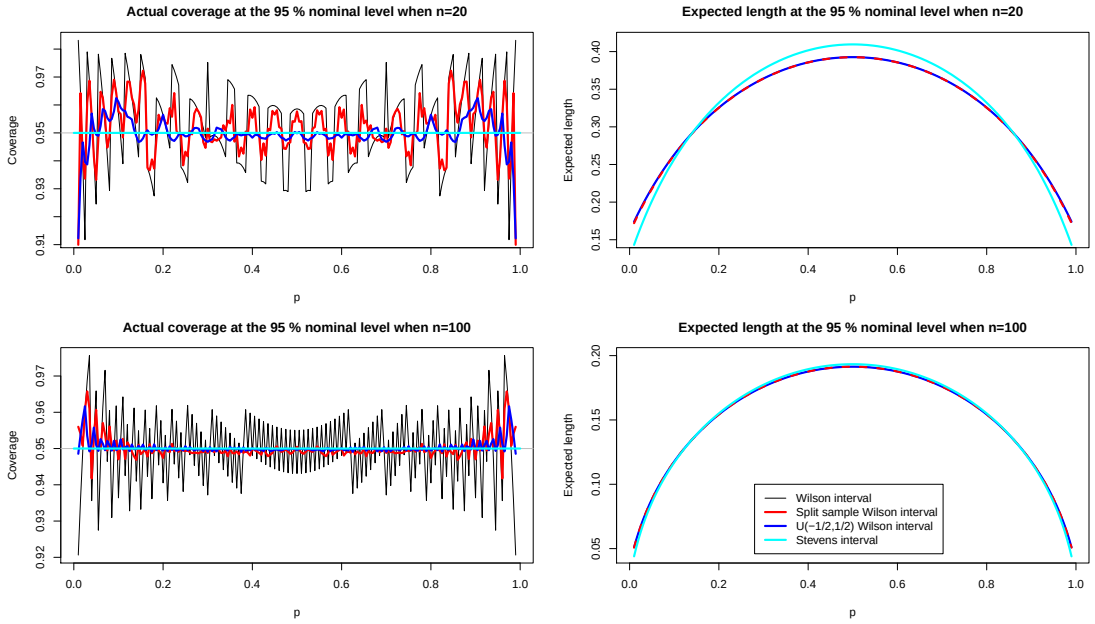
Actual coverage at the 95 % nominal level when n=20
Expected length at the 95 % nominal level when n=20
Actual coverage at the 95 % nominal level when n=100
Expected length at the 95 % nominal level when n=100
Coverage
Expected length
p
Wilson interval
Split sample Wilson interval
U(−1/2,1/2) Wilson interval
Stevens interval

and the upper bound p_L is such that

$$\nu_2 \cdot \binom{n}{X} p_U^X (1-p_U)^{n-X} + \sum_{k=0}^{X-1} \binom{n}{k} p_U^k (1-p_U)^{n-k} = \alpha/2.$$

For $\nu_1 = \nu_2 = 1$ this is the non-randomized Clopper–Pearson interval (Clopper & Pearson, 1934) and for $\nu_1 = \nu_2 = 1/2$ this is the non-randomized mid-p interval (Lancaster, 1961). The choice $\nu_1 \sim \mathrm{U}(0,1)$ and $\nu_2 = 1 - \nu_1$ corresponds to the randomized Stevens interval, which is an inversion of the most powerful unbiased test for p (Blank, 1956). We note that Stevens' procedure also can be applied when constructing confidence intervals for parameters of Poisson and negative binomial distributions.

2.2 Comparison of externally randomized intervals

The coverage fluctuations and expected length of the split sample Wilson, $\mathrm{U}(-1/2, 1/2)$ Wilson and Stevens intervals are compared in Figure 3. The $\mathrm{U}(-1/2, 1/2)$ interval improves upon the split sample interval about as much as the split sample interval improves upon the standard Wilson interval. Unlike the split sample and $\mathrm{U}(-1/2, 1/2)$ intervals, the Stevens interval has coverage exactly equal to $1-\alpha$ for

Figure 4: Range of the upper bounds of randomized confidence intervals for different combinations of n and X. The ranges are only defined for integer X; the lines are used only to guide the eye.

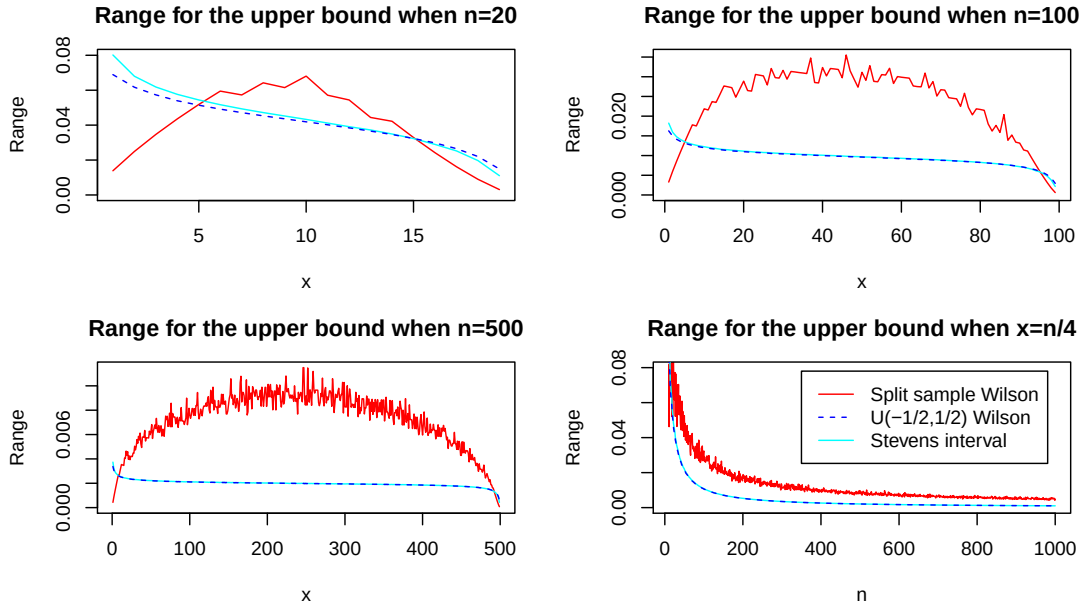

all values of p, so that there are no coverage oscillations at all. The difference in expected length between the intervals is quite small for all combinations of p and n and is decreasing in n. The Stevens interval is shorter when p is close to 0 or 1, but wider when p is close to 1/2. When $n = 20$ the difference in expected length is at most 0.018 and when $n = 100$ it is at most 0.002.

The randomization has a strong impact on the bounds of the confidence intervals. As an example, when $n = 10$ and $X = 2$ the upper bound of the split sample Wilson interval ranges from 0.476 to 0.558, meaning that the range of the upper bound is 0.082 or 18 % of the interval's conditional expected length 0.45. The impact of the randomization can be substantial even for relatively large n. Consider for instance the case $n = 200$, $X = 40$. The conditional mean length of the split sample Wilson interval is 0.11, whereas the range of the upper bound of the interval is 0.016, or 14.5 % of the conditional expected length of the interval.

The $U(-1/2, 1/2)$ Wilson and Stevens intervals are similarly affected by the randomization for small n: when $n = 10$ and $X = 2$ the conditional expected length of the Stevens interval is 0.48 and the range of the upper bound is 0.111 or 23 % of the conditional expected length. However, the impact of the randomization on these interval decreases quicker than it does for the split sample Wilson interval.

When $n = 200$ and $X = 40$ the conditional expected length of the Stevens interval is 0.11. The range of the upper bound is 0.005, which is no more than 4.5 % of the conditional expected length.

Figure 4 shows the range of the upper bounds of the three intervals for different combinations of X and n. The intervals are about as bad for $n = 20$, but the $U(-1/2, 1/2)$ Wilson and Stevens intervals are much less affected by the randomization for $n = 100$ and $n = 500$. Because of the equivariance of the intervals, the figures for the lower bound are identical if X is replaced by $n - X$.

3 Data-randomized intervals

3.1 Modifying externally randomized intervals

As pointed out by Decrouez & Hall (2014), if the result of each of the n independent trials are recorded, one can put the first n_1 observations in the first subsample and the other n_2 observations in the second subsample, so that the randomness only comes from the data itself. This means that there is no ambiguity about the interval, so that there is no need to discuss the range of the upper bound and that two different statisticians always will obtain the same confidence interval when they apply the procedure to the same data.

Alternatives to the split sample interval are available also if the sequence of observations has been recorded. It is possible to let the randomization of the Stevens interval be determined by $X_1, \ldots, X_n$, for instance by letting the sequence determine the (now approximately) $\mathrm{U}(0, 1)$ random variable ν_1. An example of this is given by the following algorithm. It generalizes the data-randomized Korn interval (Korn, 1987), for which the p-value of a rank test of the hypothesis that the 1's are uniformly distributed in the sequence is used to determine U_k. It is our hope that the below description of the procedure is somewhat more straightforward.

Algorithm: Data-randomization of the Stevens interval.

1. Compute all $\binom{n}{X}$ distinct permutations of $X_1, \ldots, X_n$.

2. Enumerate the permutations $k = 1, \ldots, \binom{n}{X}$ and associate with each permutation a number $U_k = k/\binom{n}{X}$.
   Remark. The enumeration can for instance be done by interpreting $X_1, \ldots, X_n$ as the base 2 decimal expansion of a number between 0 and 1, sorting the sequences according to this decimal expansion. The smallest U_k is then assigned to the sequence corresponding to the smallest number, the second smallest U_k to the sequence corresponding to the second smallest number, and so on.

3. For the k^* corresponding to the observed sequence $X_1, \ldots, X_n$, compute the Stevens interval with $\nu_1 = U_{k^*}$ and $\nu_2 = 1 - U_{k^*}$.

The algorithm can be applied to the U$(-1/2, 1/2)$ Wilson interval as well, with $Y = U_k - 1/2$.

The function used to associate a number U_k with each permutation is arbitrary. However, as long as this function is given explicitly, no ambiguity will remain when the above algorithm is used. We also note that the split sample function proposed by Decrouez & Hall (2014) is equally arbitrary, as is the choice of n_1 and n_2.

Data-randomized Stevens intervals can be constructed for parameters of other stochastically increasing lattice distributions, including the Poisson and negative binomial distributions, in an analogue manner. Similarly, the data-randomized U$(-1/2, 1/2)$ procedure can be used to lessen the coverage fluctuations of other confidence intervals and for other distributions.

3.2 Comparison of data-randomized intervals

When $n = 20$ the split sample method smoothens the binomial distribution by replacing its 21 possible values by 119 possible values. When the above algorithm is used to produce the Korn interval, $\sum_{k=0}^{20} \binom{20}{k} = 1,048,576$ possible values are used instead, smoothing the distribution even further, and reducing the coverage oscillations to a much greater extent.

The split sample Wilson interval is compared to the Korn and U$(-1/2, 1/2)$ Wilson intervals in Figure 5. When $n = 20$, the Korn interval has near-perfect coverage for $p \in (0.2, 0.8)$, but like the externally randomized Stevens interval it is slightly wider than the competing intervals in this region of the parameter space. This difference in expected length quickly becomes smaller as n increases. The data-randomized U$(-1/2, 1/2)$ Wilson interval has smaller coverage fluctuations than the split sample Wilson interval, but is no wider.

Figure 5: Coverage and expected length of three data-randomized confidence intervals, compared to the non-randomized Wilson interval when $n = 20$. The expected length of the Wilson interval equals that of the split sample Wilson interval and the data-randomized $\mathrm{U}(-1/2, 1/2)$ interval.

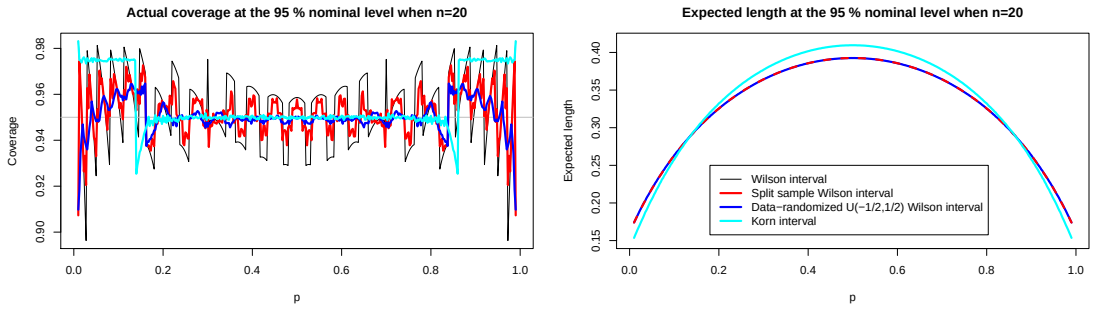

4 Discussion

In the split sample method random noise with a discrete distribution is added to the data to reduce the coverage fluctuations of a confidence interval. We have shown that it is better to use a continuous distribution such as the $\mathrm{U}(-1/2, 1/2)$ distribution for the random noise, as this lowers the coverage fluctuations even further and reduces the impact of the randomization on the interval bounds. This kind of randomization can however not be used to remove the coverage fluctuations completely. In contrast, the Stevens interval offers exact coverage. Moreover, for larger n the randomization has a smaller impact on the bounds of the Stevens interval than it has on the split sample Wilson interval. If a randomized confidence interval for the binomial proportion p is to be used, we therefore recommend the Stevens interval over split sample intervals.

The main objection against randomized confidence intervals is that different statisticians may produce different intervals when applying the same method to the same data. Decrouez & Hall (2014) solve this problem by putting the first n_1 observations in the first subsample and the remaining n_2 observations in the second subsample, using so-called data-randomization. We compared several data-randomized intervals in Section 3.2 and found that the split sample Wilson interval had the worst performance of the lot. Neither of these intervals manages to remove the coverage fluctuations completely.

When using the data-randomization strategy we have to condition our inference on an auxiliary random variable – the order of the observations – which has no relation to the parameter of interest. We then lose the rather natural property of invariance under permutations of the sample. Consequently, a statistician that reads the sequence from the left to the right may end up with a different interval

than a statistician that reads the sequence from the right to the left. Care should be taken when using data-randomization: in most binomial experiments it is arguably neither reasonable nor desirable that the order in which the observations were obtained affects the inference.

Objections against randomized confidence intervals are sometimes met by the argument that randomization already is widely used in virtually all areas of statistics: bootstrap methods, MCMC, tests based on simulated null distributions and random subspace methods (e.g. Thulin (2014)) all rely on randomization. An important difference is however that for any n the impact of the randomization on these methods can be made arbitrarily small by allowing more time for the computations. This is not possible for randomized confidence intervals for parameters of lattice distributions, where for each n there exists a positive lower bound for the impact of the randomization. In this context we also mention Geyer & Meeden (2005), who proposed a completely different approach to dealing with the ambiguity of randomized confidence intervals, based on fuzzy set theory.

In our study we applied the split sample method to the Wilson interval for a binomial proportion. This interval has excellent performance, with competitive coverage and length properties, and is often recommended for general use (Brown et al., 2001; Newcombe, 2012). We note however that it can be more or less uniformly improved upon by using a level-adjusted Clopper–Pearson interval (Thulin, 2013b) instead; we opted to use the Wilson interval in our study simply because it is less computationally complex.

Acknowledgement

The author wishes to thank Silvelyn Zwanzig for helpful discussions.

References

Blank, A.A. (1956). Existence and uniqueness of a uniformly most powerful randomized unbiased test for the binomial. *Biometrika*, **43**, 465-466.

Brown, L.D., Cai, T.T., DasGupta, A. (2001). Interval estimation for a binomial proportion. *Statistical Science*, **16**, 101–133.

Clopper, C.J., Pearson, E.S. (1934). The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*, **26**, 404–413.

Decrouez, G., Hall, P. (2013). Normal approximation and smoothness for sums of means of lattice-valued random variables. *Bernoulli*, **19**, 1268–1293.

Decrouez, G., Hall, P. (2014). Split sample methods for constructing confidence intervals for binomial and Poisson parameters. *Journal of the Royal Statistical Society: Series B*, DOI: 10.1111/rssb.12051.

Geyer, C.J., Meeden, G.D. (2005). Fuzzy and randomized confidence intervals and P-values. *Statistical Science*, **20**, 358–366.

Göb, R. Lurz, K. (2013). Design and analysis of shortest two-sided confidence intervals for a probability under prior information. *Metrika*, DOI: 10.1007/s00184-013-0445-9.

Gonçalves, L., de Oliviera, M.R., Pascoal, C., Pires, A. (2012). Sample size for estimating a binomial proportion: comparison of different methods. *Journal of Applied Statistics*, **39**, 2453–2473.

Korn, E.L. (1987). Data-randomized confidence intervals for discrete distributions. *Communications in Statistics: Theory and Methods*, **16**, 705–715.

Krishnamoorthy, K., Peng, J. (2011). Improved closed-form prediction intervals for binomial and Poisson distribution. *Journal of Statistical Planning and Inference*, **141**, 1709–1718.

Lancaster, H.O. (1961). Significance tests in discrete distributions. *Journal of the American Statistical Association*, **56**, 223–234.

Newcombe, R.G. (2012). *Confidence Intervals for Proportions and Related Measures of Effect Size*. Chapman & Hall.

Newcombe, R.G. (2011). Measures of location for confidence intervals for proportions. *Communications in Statistics: Theory and Methods*, **40**, 1743–1767.

Stevens, W. (1950). Fiducial limits of the parameter of a discontinuous distribution. *Biometrika*, **37**, 117-129.

Thulin, M. (2013a). The cost of using exact confidence intervals for a binomial proportion. arXiv:1303.1288.

Thulin, M. (2013b). Coverage-adjusted confidence intervals for a binomial proportion. *Scandinavian Journal of Statistics*, DOI: 10.1111/sjos.12021.

Thulin, M. (2014). A high-dimensional two-sample test for the mean using random subspaces. *Computational Statistics & Data Analysis*, **74**, 26–38.

Wilson, E.B. (1927). Probable inference, the law of succession and statistical inference. *Journal of the American Statistical Association*, **22**, 209–212.

Contact information:
Måns Thulin, Department of Mathematics, Uppsala University, Box 480, 751 06 Uppsala, Sweden. E-mail: thulin@math.uu.se